

Engineering Rules

This section provides engineering rules or guidelines that must be considered prior to configuring the system for your network deployment.

- APN Engineering Rules, on page 1
- DHCP Service Engineering Rules, on page 1
- GGSN Engineering Rules, on page 1
- GRE Tunnel Interface and VRF Engineering Rules, on page 2
- GTP Engineering Rules, on page 2
- Interface and Port Engineering Rules, on page 2
- Lawful Intercept Engineering Rules, on page 4
- MBMS Bearer Service Engineering Rules, on page 4
- Service Engineering Rules, on page 4
- Subscriber Engineering Rules, on page 5

APN Engineering Rules

The following engineering rules apply to APNs:

- APNs must be configured within the context used for authentication.
- A maximum of 2,048 APNs per system can be configured.

DHCP Service Engineering Rules

The following engineering rule applies to the DHCP Service:

- Up to 8 DHCP servers may be configured per DHCP service.
- A maximum of 3 DHCP server can be tried for a call.

GGSN Engineering Rules

The following engineering rules apply when the system is configured as a GGSN:

- Gn/Gp interfaces can be configured. That is, if a system context is configured with a GGSN service, then all interfaces in that context may be used.
- Gi interfaces can be configured. That is, if a system context is configured as a destination context for an APN, then all interfaces in that context may be used.
- Ga interfaces. That is, if a system context is configured for GTPP accounting, then all interfaces in that context may be used.
- One GSN-MAP node may be configured per system context (in lieu of Gc).
- Up to 1000 network requested PDP contexts may be configured.
- Up to 8 GTPP groups (excluding the default GTPP group) can be configured per chassis.
- Up to 4 GTPP Storage Servers can be configured per GTPP group.
- Up to 32 GTPP Storage Servers can be configured per system context.
- Up to 511 GRE tunnel interface can be configured per context.

GRE Tunnel Interface and VRF Engineering Rules

The following engineering rules apply to GRE tunnel interface and VRF contexts:

- A maximum of 511 GRE tunnels are allowed to configure in a context but subject to maximum of 2048 GRE tunnels per chassis.
- A maximum of 300 virtual routing and forwarding (VRF) tables are allowed to be configured in a context, subject to a maximum of 2,048 VRFs per chassis.
- A maximum of 10000 IP routes in Release 9.0 and 16384 IP routes in Release 10.0 onward are supported in a VRF context configuration mode.

GTP Engineering Rules

The following engineering rules apply to GTP on GGSN:

- A maximum of 11 primary (no secondary) PDP context per subscriber can be configured.
- A maximum of 1 primary and 10 secondary PDP context per subscriber can be configured.

Interface and Port Engineering Rules

The rules discussed in this section pertain to both the Ethernet 10/100 and Ethernet 1000 Line Cards and the four-port Quad Gig-E Line Card (QGLC) and the type of interfaces they facilitate.

Pi Interface Rules

This section describes the engineering rules for the Pi interface.

FA to HA Rules

When supporting Mobile IP, the system can be configured to perform the role of an FA, an HA, or both. This section describes the engineering rules for the Pi interface when using the system as a FA.

The following engineering rules apply to the Pi interface between the FA and HA:

- A Pi interface is created once the IP address of a logical interface is bound to an FA service.
- The logical interface(s) that will be used to facilitate the Pi interface(s) must be configured within the egress context.
- FA services must be configured within the egress context.
- If the system is configured as a FA is communicating with a system configured as an HA, then it is recommended that the name of the context in which the FA service is configured is identical to the name of the context that the HA service is configured in on the other system.
- Each FA service may be configured with the Security Parameter Index (SPI) of the HA that it will be communicating with over the Pi interface.
- Multiple SPIs can be configured within the FA service to allow communications with multiple HAs over the Pi interface. It is best to define SPIs using a netmask to specify a range of addresses rather than entering separate SPIs. This assumes that the network is physically designed to allow this communication.
- Depending on the services offered to the subscriber, the number of sessions facilitated by the Pi interface can be limited.

HA to FA

The following engineering rules apply to the Pi interface between the HA and FA:

- When supporting Mobile IP, the system can be configured to perform the role of a FA, an HA or both. This section describes the engineering rules for the Pi interface when using the system as an HA.
- A Pi interface is created once the IP address of a logical interface is bound to an HA service.
- The logical interface(s) that will be used to facilitate the Pi interface(s) must be configured within an ingress context.
- HA services must be configured within an ingress context.
- If the system configured as an HA is communicating with a system configured as a FA, then it is recommended that the name of the context in which the HA service is configured is identical to the name of the context that the FA service is configured in on the other system.
- Each HA service may be configured with the Security Parameter Index (SPI) of the FA that it will be communicating with over the Pi interface.
- Multiple SPIs can be configured within the HA service to allow communications with multiple FAs over the Pi interface. It is best to define SPIs using a netmask to specify a range of addresses rather than entering separate SPIs. This assumes that the network is physically designed to allow this communication.
- Each HA service must be configured with a Security Parameter Index (SPI) that it will share with mobile nodes.
- Depending on the services offered to the subscriber, the number of sessions facilitated by the Pi interface can be limited in order to allow higher bandwidth per subscriber.

GRE Tunnel Interface Rule

The following engineering rules apply to the GRE tunnel interface between two GRE tunnel nodes:

- A maximum of 512 IP tunnels (511 GRE tunnels + 1 not tunnel interfaces) are allowed to configure in a context but subject to a maximum of 2048 GRE tunnels per chassis.

Lawful Intercept Engineering Rules

The following engineering rules apply to Lawful Intercept on supported AGW service:

- A maximum of 1000 Lawful Intercepts can be performed simultaneously.

MBMS Bearer Service Engineering Rules

The following engineering rules apply to MBMS bearer services:

- A maximum 15 downlink SGSNs per MBMS bearer service are supported on ST16.
- A maximum 225 downlink SGSNs per MBMS bearer service are supported on the system.
- A maximum of 2 BMSC (1 primary and 1 secondary) supported per MBMS bearer service.

Service Engineering Rules

The following engineering rules apply to services configured within the system:

- A maximum of 256 services (regardless of type) can be configured per system.



Caution Large numbers of services greatly increase the complexity of management and may impact overall system performance (i.e. resulting from such things as system handoffs). Therefore, it is recommended that a large number of services only be configured if your application absolutely requires it. Please contact your local service representative for more information.

- Up to 2,048 MN-HA and 2048 FA-HA SPIs can be supported for a single HA service.
- Up to 2,048 FA-HA SPIs can be supported for a single FA service.
- The system supports unlimited peer FA addresses per HA.
- The system maintains statistics for a maximum of 8192 peer FAs per HA service.
- If more than 8192 FAs are attached, older statistics are identified and overwritten.
- The system maintains statistics for a maximum of 4096 peer HAs per FA service.
- The total number of entries per table and per chassis is limited to 256.

- Up to 10,000 LAC addresses can be configured per LNS service.

Caution Even though service names can be identical to those configured in different contexts on the same system, this is not a good practice. Having services with the same name can lead to confusion, difficulty in troubleshooting the problems, and make it difficult to understand outputs of **show** commands.

Subscriber Engineering Rules

The following engineering rule applies to subscribers configured within the service:

- Default subscriber templates may be configured on a per FA service basis.